

Search Results

Search Results for: [prefix <near> match<AND>((node <near> leaf<AND>((direct <near> table<AND>((full <near> match<AND>((longest <near> match)))))))]

Found 18 of 96,434 searched. → Rerun within the Portal

Search within Results

 > Advanced Search > Search Help/Tips

Sort by: Title Publication Publication Date Score  Binder

Results 1 - 18 of 18 short listing

1	Scalable high-speed prefix matching	95%
	Marcel Waldvogel , George Varghese , Jon Turner , Bernhard Plattner ACM Transactions on Computer Systems (TOCS) November 2001 Volume 19 Issue 4 Finding the longest matching prefix from a database of keywords is an old problem with a number of applications, ranging from dictionary searches to advanced memory management to computational geometry. But perhaps today's most frequent best matching prefix lookups occur in the Internet, when forwarding packets from router to router. Internet traffic volume and link speeds are rapidly increasing; at the same time, a growing user population is increasing the size of routing tables against which p ...	
2	Pattern Matching in Trees	87%
	Christoph M. Hoffmann , Michael J. O'Donnell Journal of the ACM (JACM) January 1982 Volume 29 Issue 1	
3	Small forwarding tables for fast routing lookups	87%
	Mikael Degermark , Andrej Brodnik , Svante Carlsson , Stephen Pink ACM SIGCOMM Computer Communication Review , Proceedings of the ACM SIGCOMM '97 conference on Applications, technologies, architectures, and protocols for computer communication October	

4 Routing on longest-matching prefixes 37%
 Willibald Doeringer , Günter Karjoth , Mehdi Nassehi
IEEE/ACM Transactions on Networking (TON) February 1996
Volume 4 Issue 1

5 External memory algorithms and data structures 36%
 Jeffrey Scott Vitter
ACM Computing Surveys (CSUR) June 2001
Volume 33 Issue 2
Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

6 Router plugins 22%
 Dan Decasper , Zubin Dittia , Guru Parulkar , Bernhard Plattner
IEEE/ACM Transactions on Networking (TON) February 2000
Volume 8 Issue 1

7 A guided tour to approximate string matching 20%
 Gonzalo Navarro
ACM Computing Surveys (CSUR) March 2001
Volume 33 Issue 1
We survey the current techniques to cope with the problem of string matching that allows errors. This is becoming a more and more relevant issue for many fast growing areas such as information retrieval and computational biology. We focus on online searching and mostly on edit distance, explaining the problem and its relevance, its statistical behavior, its history and current developments, and the central ideas of the algorithms and their complexities. We present a number of experiments to ...

8 Optimal amnesic probabilistic automata or how to learn and 12%
 classify proteins in linear time and space
Alberto Apostolico , Gill Bejerano
Proceedings of the fourth annual international conference on Computational molecular biology April 2000

Statistical modeling of sequences is a central paradigm of machine learning that finds multiple uses in computational molecular biology and many other domains. The probabilistic automata typically built in these contexts are subtended by uniform, fixed-memory Markov models. In practice, such automata tend to be unnecessarily bulky and computationally imposing both during their synthesis and use. In [8], much more compact, tree-shaped variants of probabilistic automata are built which assume ...

9 Data compression with finite windows

11%

 E. R. Fiala , D. H. Greene

Communications of the ACM April 1989

Volume 32 Issue 4

Several methods are presented for adaptive, invertible data compression in the style of Lempel's and Ziv's first textual substitution proposal. For the first two methods, the article describes modifications of McCreight's suffix tree data structure that support cyclic maintenance of a window on the most recent source characters. A percolating update is used to keep node positions within the window, and the updating process is shown to have constant amortized cost. Other methods explore the ...

10 Algorithms for trie compaction

7%

 M. Al-Suwaiyel , E Horowitz

ACM Transactions on Database Systems (TODS) May 1984

Volume 9 Issue 2

The trie data structure has many properties which make it especially attractive for representing large files of data. These properties include fast retrieval time, quick unsuccessful search determination, and finding the longest match to a given identifier. The main drawback is the space requirement. In this paper the concept of trie compaction is formalized. An exact algorithm for optimal trie compaction and three algorithms for approximate trie compaction are given, and an analysis of the ...

11 Packet classification using tuple space search

6%

 V. Srinivasan , S. Suri , G. Varghese

ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication August 1999

Volume 29 Issue 4

12 Router plugins: a software architecture for next generation routers

6%



Dan Decasper , Zubin Dittia , Guru Parulkar , Bernhard Plattner

ACM SIGCOMM Computer Communication Review, Proceedings of the ACM SIGCOMM '98 conference on Applications, technologies, architectures, and protocols for computer communication October 1998
Volume 28 Issue 4

13 Large-scale assembly of DNA strings and space-efficient 2%
 ④ construction of suffix trees
 S. Rao Kosaraju , Arthur L. Delcher
 Proceedings of the twenty-seventh annual ACM symposium on Theory of computing May 1995

14 All searches are divided into three parts 2%
 ④ David E. Siegel
 ACM SIGAPL APL Quote Quad , Proceedings of the APL98 conference on Array processing language July 1998
 Volume 29 Issue 3

15 Designing a global name service 1%
 ④ Butler W Lampson
 Proceedings of the fifth annual ACM symposium on Principles of distributed computing November 1986

16 Practical dictionary management for hardware data compression 1%
 ④ Suzanne Bunton , Gaetano Borriello
 Communications of the ACM January 1992
 Volume 35 Issue 1

17 LeZi-update 0%
 ④ Amiya Bhattacharya , Sajal K. Das
 Wireless Networks March 2002
 Volume 8 Issue 2/3
 The complexity of the mobility tracking problem in a cellular environment has been characterized under an information-theoretic framework. Shannon's entropy measure is identified as a basis for comparing user mobility models. By building and maintaining a dictionary of individual user's *path* updates (as opposed to the widely used location updates), the proposed adaptive on-line algorithm can learn subscribers' profiles. This technique evolves out of the concepts of lossless compression. T ...

18 Compressed suffix arrays and suffix trees with applications to 0%

 text indexing and string matching (extended abstract)
Roberto Grossi , Jeffrey Scott Vitter
Proceedings of the thirty-second annual ACM symposium on Theory of computing May 1999

Results 1 - 18 of 18 short listing

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2002 ACM, Inc.